

My Research

What is a concept? What does it mean to be a concept-possessing, concept-using agent?

My research involves the extension of Peter Gärdenfors' conceptual spaces theory of concepts, attempting to fill in the gaps in his theory and show how it can provide a bridge not only between associationist and symbolic accounts of cognition (as he envisions it) but also between an atomic theory of concepts qua Jerry Fodor and a non-atomic theory qua Jesse Prinz. Fodor and Prinz present two of the currently more popular and debated theories of concepts that have sprung up as replacements for the discredited classical definitionist and imagist accounts. The incommensurability of their accounts is, I believe, only apparent. Not only are both views useful, I believe they (or something like them) are both necessary.

Per Fodor, I want to emphasize the centrality of an explicit and coherent theory of concepts to any work in cognitive science or AI. Too much of the time the theory of concepts underlying work in the field has been taken for granted if not simply implicitly assumed; and my suspicions are with Fodor that that theory has in many substantial ways been wrong.

My thesis revolves around four main points. First of all, any theory of concepts, to be complete, must reconcile contrasting theoretical understandings of concepts along at least two dimensions: concepts as conceptually atomic vs. concepts as composites of other concepts, and concepts as shared entities in a public (social) space vs. concepts as private entities.

Second, a theory of concepts grounded in the metaphor of geometry, in the style of Gärdenfors' conceptual spaces theory, can provide a means of bridging these divides, between the conceptually atomic and the conceptually structured, and between the public and the private. Third, a single shape or a small set of shapes – building blocks – can provide perceptual primitives to ground conceptual spaces theory, giving common underlying form to all the concepts one has need to talk about. Fourth, a theory of concepts is best tested in some form of implementation that can then be used in refining the theory, creating a tight theory-implicit model-explicit model-implementation-theory loop.

Its Relationship to Enaction

Understanding cognition is more than recognizing the role of context: many AI researchers twenty years ago, when I was doing my undergraduate course, understood the importance of context. It's not just that a cognitive agent is embedded in a certain environment or embodied in a certain way that constrains his cognition; it's that cognition is largely a process of dynamic, physical engagement with the environment, in which the line separating organism from environment may from some important perspectives appear quite arbitrary.

My supervisor Ron Chrisley and I have, as part of a separate project, been exploring a way of approaching enactive experience that preserves a form of representationalism – one in which the contents of current experience go beyond current sensory input. This might be called counterfactualist representationalism: if I looked over there *this* is what I'd expect to see, because e.g. that's what I saw the last time I looked there. In short,

experience is more than current sensory experience, and current sensory experience is more than current sensory input. The medium for our research has been a robotic dog that assembles a composite depiction of its visual “experience” one jigsaw-puzzle-piece saccade at a time.

Noë’s enactive approach is more strictly forward looking than the one I’ve been taking with my supervisor, in which present experience is largely constrained by past experience: everything that has led up to the present moment. Much of the time one sees what past experience has taught one to expect to see.

Although our work to date has focused entirely on specifying the non-conceptual contents of experience, I believe that some of the same methods can be used for specifying the conceptual contents as well. Indeed I think that all experience is conceptualized to greater or lesser degree – by which I mean meeting some form of Gareth Evans’ Generality Constraint – and it is hard to conceive of experience that is fully non-conceptualized. As abstract thinking is built upon dynamic physical engagement with the world, as *knowing how* underlies much if not all of *knowing that*, so, too, is conceptual knowledge and experience built upon the non-conceptual: concepts projected over top of non-conceptualized experience and all but obscuring it, so that we might even think sometimes that the non-conceptualized experience does not exist.

Conceptual knowledge is not a collection of dictionary-style definitions nor is it a *process* of collecting such definitions. Though concepts may *look* like definitions when we try to explain them, they are nonetheless the result of our dynamic engagement both with our environment as a whole and with the society of which we are members: so I have my personal concept DOG, which may vary in greater or lesser ways from the next person, or from my own concept DOG at different points in time; and I have the concept DOG that is part of the social space in which we all share.

My Question

The question I would like to pose to the working group is this: despite all the technical and theoretical advances that have been made, it is still difficult to construct an artefact that is engaged with its environment in the rich sort of way required by an enactive approach¹ – never mind one that qualifies as a concept-possessing, concept-using cognitive agent. Nonetheless, one might well wish to construct artefacts with which to better understand cognition and conceptualization. Question: how much of the cognitive and enactive requirements might be put off onto a user (e.g., a test subject) dynamically engaged with the artefact (which might, then, be something not so different from a traditional computer program)?

References

Chrisley, R. & Parthemore, J. (2007) *Synthetic phenomenology: Exploiting embodiment to specify the non-conceptual content of visual experience*, *Journal of Consciousness Studies* 14(7): pp. 44-58.

¹ Indeed, many of the GOFAI projects of old may have been so context-impooverished simply because people built what they most easily could with the tools available.

Fodor, J. (1998) Concepts: Where Cognitive Science Went Wrong. Clarendon Press, Oxford.

Gärdenfors, P. (2004) Conceptual Spaces: The Geometry of Thought, Cambridge, USA: Bradford Books.

Prinz, J. (2004) Furnishing the Mind: Concepts and their Perceptual Basis, Cambridge, USA: MIT Press.